Joint Power and Resource Block Allocation for Delay-Reliable Downlink XR Scheduling in 6G

Moyukh Laha*, Goutam Das†, Gabriel-Miro Muntean*
*School of Electronic Engineering, Dublin City University, Ireland
†Dept. of Electronics & Electrical Communication Engineering, Indian Institute of Technology Kharagpur, India
Email: {moyukh.laha, gabriel.muntean}@dcu.ie, gdas@ece.iitkgp.ac.in

***Abstract*—Emerging 6G networks position eXtended Reality (XR) as a key service class that simultaneously requires very high throughput and stringent latency and reliability guarantees. Existing delay- and reliability-aware schedulers optimize resource block (RB) allocation under fixed transmit power and therefore ignore the intrinsic coupling between transmit power and achievable rate. In this work, we propose Power-Optimized Delay-Reliable XR Scheduling (PODRX), a joint power and RB allocation framework that leverages the PDU Set Delay Tracking Queue (PDTQ) mechanism for PDU Set delay budget (PSDB) monitoring. The scheduling problem is formulated as a Mixed-Integer Nonlinear Program (MINLP) and systematically transformed into a tractable Mixed-Integer Linear Program (MILP), which enables computation of globally optimal solutions. Simulation results show that PODRX reduces transmit power consumption by more than 50% compared with state-of-the-art fixed-power delay-reliability-aware XR schedulers, while simultaneously increasing the satisfied XR user count by approximately 20%.**

***Index Terms*—Extended reality, PDU Set, joint power-RB optimization, delay-reliable scheduling, 6G, XR.**

I. Introduction

Extended Reality (XR), which encompasses Virtual Reality (VR), Augmented Reality (AR), and Mixed Reality (MR), is transforming how users interact with machines and with one another. The 6G roadmaps identify XR as a flagship service class [1] that imposes a distinctive combination of requirements: sustained high throughput for volumetric content delivery, sub-millisecond-scale latency, and stringent reliability to preserve perceptual immersion. Taken together, these requirements position XR at the intersection of *enhanced Mobile Broadband* (eMBB) and *Ultra-Reliable Low-Latency Communication* (URLLC) [1], a combined service regime that remains largely unaddressed by existing 5G designs.

XR traffic has an inherently frame-oriented structure. At the application layer, each rendered frame is segmented into multiple *Protocol Data Units* (PDUs), all of which must arrive at the receiver before a common deadline to enable correct decoding. To capture this interdependence, 3GPP Release 18 introduced the PDU Set abstraction [2], [3], which groups all PDUs belonging to the same media frame and associates them with a *PDU Set Delay Budget* (PSDB). If even one PDU reaches its destination later than the PSDB, the entire PDU Set becomes invalid and the resources spent on transmitting these PDUs are wasted [4]. Furthermore, 3GPP redefined cellular XR capacity as the number of XR users that can be simultaneously supported [5], where a user is deemed *satisfied* only if a sufficiently high fraction (typically 99% under the 3GPP specification) of its PDU Sets meet their PSDB. In this work, we refer to this joint requirement on delay and reliability as *delay-reliability*.

Designing MAC-layer schedulers that explicitly enforce delay-reliability targets remains challenging. Classical methods such as *Proportional Fair* (PF) treat packets independently and are therefore ill-suited to PSDB-constrained XR flows. Several recent proposals seek to narrow this gap. The work in [6] formulates scheduling as a periodic Markov Decision Process but assumes uniform packet sizes, [7] considers frame-level transmissions without providing explicit delay-reliability guarantees, and [4] incorporates deadline awareness yet does not include an explicit delay tracking mechanism. Architectural enhancements, such as the application-aware MAC in [8] and AI-assisted provisioning in [9], improve user satisfaction but do not target PDU Set-level XR scheduling.

In response to these limitations, the Delay-Reliability Optimal XR (DROX) scheduler [10] introduced the Delay Tracking Queue (DTQ), a virtual queueing construct that maintains per-user, per-deadline queues and allows the scheduler to capture the urgency of each pending Data Unit (DU). Subsequent work extended DTQs to multisensory immersive XR services [11]. However, all DTQ-based schedulers developed so far treat transmit power as a fixed system parameter and optimize only over the resource block (RB) domain. As a result, they overlook a fundamental degree of freedom: the coupling between transmit power and achievable data rate. Increasing the power allocated to a given user can raise its per-RB throughput, which potentially reduces the number of RBs required and frees spectrum for other users, or enables the same throughput to be achieved at substantially lower energy cost. Although joint power and RB allocation has been extensively studied in conventional wireless systems, it has not yet been integrated into a delay-reliability-aware framework tailored to XR. This paper addresses this gap by introducing PODRX (Power-Optimized Delay-Reliable XR Scheduling).

The main contributions are summarized as follows:

- We propose a PDU Set Delay Tracking Queue (PDTQ) mechanism for PSDB tracking that supports the specific requirements of XR applications.

- Using the PDTQ structure, we formulate the joint power and RB allocation problem for a delay-reliability-aware XR scheduler as a Mixed-Integer Nonlinear Program (MINLP) that maximizes the number of delay-reliability-satisfied XR users subject to per-slot power and RB constraints.
- To solve the MINLP, we systematically derive an equivalent Mixed-Integer Linear Program (MILP) that can be solved to global optimality with standard optimization solvers.
- Through extensive simulations, we show that PODRX provides significant gains compared with fixed-power delay-reliability-aware schedulers and multifold improvements compared with traditional schedulers.

The remainder of this paper is organized as follows. Section II describes the system model and the PDTQ mechanism. Section III presents the PODRX formulation and its MILP transformation. Section IV reports the performance analysis. Finally, Section V concludes the paper.

II. PDU Set Delay Budget Tracking Mechanism

A. System Model and Assumptions

We consider a single gNB operating under the standard 5G/6G protocol stack and serving a set of XR users. Application-layer video frames are organized as PDU Sets that are buffered in per-user Medium Access Control queues (MACQs) at the gNB, each associated with a PDU Set Delay Budget (PSDB). The scheduler operates at the MAC layer on the downlink only; uplink scheduling is outside the present scope. We focus on video-centric XR traffic, which dominates bandwidth consumption due to volumetric rendering. Each user is assumed to run a single XR application within a dedicated RAN slice allocated exclusively to XR services, consistent with the slice-based isolation envisioned in 3GPP Release 18 [12]. Coexistence with non-XR traffic classes is deferred to future work. Moreover, in order to solve the problem optimally, we assume full knowledge of future PDU Set arrivals and channel realizations over the entire scheduling horizon. This idealized assumption is infeasible for real-time operation but is essential to establish a theoretical upper bound on achievable performance. Relaxing this assumption with a more practical approach constitutes a natural next step, which we leave for future work.

B. PDTQ Structure and Operation

Building on the delay tracking mechanism in [10], we introduce a key enhancement, namely a PDU Set-level delay tracking mechanism. The proposed mechanism consists of *PDU Set Delay Tracking Queues (PDTQs)* and a configurable *Delay Tracking Granularity (DTG)*.

PDTQ architecture: For each user i, we maintain K_i PDTQs, numbered from K_i (containing the freshest arrivals, farthest from their PSDB expiry) down to 1 (containing the most time-critical PDU Sets). A sentinel index 0 collects PDU Sets whose PSDBs have already elapsed, that is, delay violations.

Table I
Symbol Notation

Symbol	Description
N	Number of XR users
T	Scheduling horizon (time slots)
C	Total available resource blocks per slot
P_{tot}	Base station aggregate power budget (W) per slot
L	Number of candidate power levels
P_l	Transmit power at level l (W)
$Q_i^j(t)$	PDU Set in PDTQ-j for user i at slot t
K_i	Highest PDTQ index for user i
$\mathcal{J}_i$	Index set $\{1, 2, \ldots, K_i\}$
μ_i	Maximum tolerable violation ratio for user i
$\gamma_i(T)$	Measured delay-reliability violation ratio
$a_i(t)$	Binary arrival indicator at slot t
$f_i(t)$	Size (bits) of the PDU Set arriving at slot t
$R_{i,l}(t)$	Per-RB throughput at power level l
$x_i^j(t)$	RBs assigned to PDTQ-j of user i
$z_i^l(t)$	Binary power level selector
$\mathbb{I}\{\cdot\}$	Indicator function

The temporal resolution of this indexing is governed by the DTG τ, which we set equal to the Transmission Time Interval (TTI). A newly arrived PDU Set with residual PSDB Δ_i is placed into the PDTQ at index $\phi = \lfloor(\Delta_i - \tau)/\tau\rfloor$.

PDTQ dynamics: In every scheduling interval, the scheduler selects PDU Sets from the PDTQs for transmission (the selection policy is detailed in Section III). Any PDU Set that is not scheduled in a given interval migrates one position closer to its deadline, that is, from PDTQ-j to PDTQ-$(j-1)$. A PDU Set that resides in PDTQ-1 at the end of an interval and still has not been served is declared a delay violation and discarded. This progressive shifting provides the scheduler with an explicit per-PDU-Set measure of temporal urgency. The PDTQ framework and delay tracking mechanism are illustrated in Figs. 1 and 2.

III. PODRX: Joint Power-RB Scheduling

Having established the PDTQ mechanism, we now formulate the PODRX joint power and RB allocation problem. Table I summarizes the principal notation. The core objective is to jointly optimize RB allocation and power assignment so that the number of delay-reliability-violated XR users is minimized under given resource constraints.

We denote the discrete power level index set by $\mathcal{L} \triangleq \{1, 2, \ldots, L\}$.

A. MINLP Formulation

Decision variables: Two interdependent sets of decisions are taken at each downlink scheduling slot t:

$$x_i^j(t) \in \{0, 1, \ldots, C\} \quad \text{(RB assignment)} \tag{1}$$

$$z_i^l(t) \in \{0, 1\} \quad \text{(power level indicator)} \tag{2}$$

Here, $x_i^j(t)$ denotes the number of RBs allocated to PDTQ-j of user i at time t, and $z_i^l(t) = 1$ indicates that user i operates at power level P_l during slot t. Based on these variables, the instantaneous transmit power is given by $p_i(t) = \sum_{l\in\mathcal{L}} P_l z_i^l(t)$, and the corresponding per-RB throughput is $\alpha_i(t) = \sum_{l\in\mathcal{L}} R_{i,l}(t) z_i^l(t)$. The per-RB throughput $R_{i,l}(t)$ is

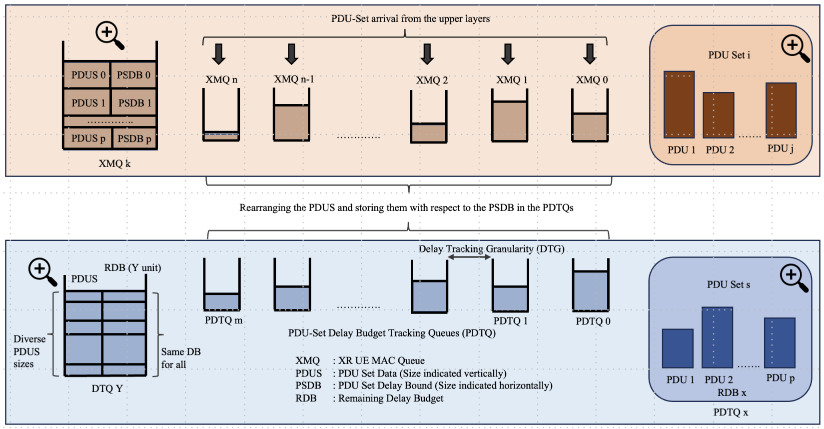

Figure 1. PDTQ mechanism: PDU sets from the upper layers come to the individual XR user MAC queue, which are then placed in the appropriate PDTQs based on their PSDB. We zoom in on a random XR MAC queue and one PDTQ to further illustrate the concept.

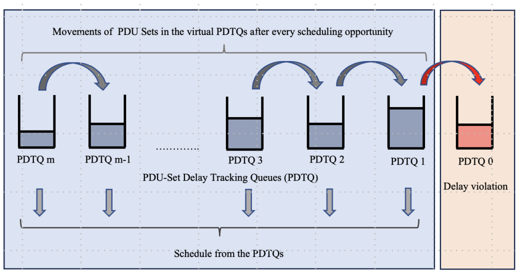

Figure 2. PDU Set movements through the PDTQs: precise delay tracking.

derived by computing the SINR from the transmit power and channel gain, mapping it to an MCS index, and obtaining the corresponding spectral efficiency.

The variable $z_i^l(t)$ is indexed per user and per slot, which grants the optimizer full flexibility to adapt each user's power level at every TTI.

1) Objective Function: The objective is to minimize the number of users whose delay-reliability performance fails to meet their respective targets:

$$O: \quad \text{minimize} \sum_{i=1}^{N} \mathbb{I}\{\gamma_i(T) \geq \mu_i\}. \tag{3}$$

2) Queue Dynamics Constraints: Arrival placement: Upon generation, each PDU Set is inserted into the highest-indexed PDTQ:

$$C1: \quad Q_i^{K_i}(t+1) = a_i(t) f_i(t), \quad \forall i, t, \tag{4}$$

where K_i maps the PSDB to the appropriate PDTQ depth.

Temporal progression and service: Unserved PDU Sets shift closer to the deadline and are reduced by the volume served by the scheduler:

$$C2: \quad Q_i^{j-1}(t+1) = Q_i^j(t) - \alpha_i(t) x_i^j(t), \quad \forall i, t, j \in \mathcal{J}_i. \tag{5}$$

3) Delay-Reliability Constraints: Violation accounting: The fraction of PDU Set arrivals that miss their PSDB characterizes delay-reliability:

$$C3: \quad \gamma_i(T) = \frac{\sum_{t=1}^{T} \mathbb{I}(Q_i^0(t) > 0)}{\sum_{t=1}^{T} a_i(t)}, \quad \forall i. \tag{6}$$

Service feasibility: The served volume from any PDTQ cannot exceed its current occupancy:

$$C4: \quad \alpha_i(t) x_i^j(t) \leq Q_i^j(t), \quad \forall i, t, j \in \mathcal{J}_i. \tag{7}$$

4) Resource Constraints: Spectrum budget: The aggregate RB allocation across all users and queues is bounded by the available bandwidth:

$$C5: \quad \sum_{i=1}^{N} \sum_{j=1}^{K_i} x_i^j(t) \leq C, \quad \forall t. \tag{8}$$

Power budget: The total transmit power must respect the base station's power budget per slot:

$$C6: \quad \sum_{i=1}^{N} \sum_{l \in \mathcal{L}} P_l z_i^l(t) \leq P_{\text{tot}}, \quad \forall t. \tag{9}$$

Exclusive power selection: Each user must be assigned exactly one power level in every slot:

$$C7: \quad \sum_{l \in \mathcal{L}} z_i^l(t) = 1, \quad \forall i, t. \tag{10}$$

5) Variable Domain Constraints:

$$C8: \quad x_i^j(t) \in \{0, 1, \ldots, C\}, \quad \forall i, t, j \in \mathcal{J}_i \tag{11}$$

$$C9: \quad z_i^l(t) \in \{0, 1\}, \quad \forall i, l, t \tag{12}$$

$$C10: \quad \alpha_i(t) \geq 0, \quad \forall i, t. \tag{13}$$

6) Complete Formulation: Collecting the above, the complete PODRX problem is

$$\boxed{\begin{aligned} O: \quad & \text{minimize} \sum_{i=1}^{N} \mathbb{I}\{\gamma_i(T) \geq \mu_i\} \\ \text{subject to:} \quad & C1 - C10. \end{aligned}} \tag{14}$$

Three sources of nonlinearity render this an MINLP: (i) the integer-valued RB variables, (ii) the indicator functions in both the objective and the violation metric, and (iii) the bilinear coupling $\alpha_i(t)x_i^j(t)$ that appears in the queue evolution constraint $C2$ and the service feasibility constraint $C4$, where the per-RB rate $\alpha_i(t)$ itself depends on the binary power selections $z_i^l(t)$. In the following, we remove these nonlinearities through a systematic reformulation.

B. Linearization

We now convert the MINLP into an equivalent MILP that can be solved by standard solvers.

1) McCormick Linearization for Bilinear Terms: Expanding the per-RB rate, the bilinear term in $C2$ and $C4$ can be written as $\sum_{l\in\mathcal{L}} R_{i,l}(t)\, z_i^l(t)\, x_i^j(t)$. The nonlinearity arises from the product $z_i^l(t)\, x_i^j(t)$ of a binary and a bounded integer variable. We capture this product exactly by introducing auxiliary variables $w_i^{j,l}(t) \triangleq x_i^j(t)\, z_i^l(t)$ together with McCormick envelope constraints:

$$C11: \quad w_i^{j,l}(t) \le C \cdot z_i^l(t) \tag{15}$$

$$C12: \quad w_i^{j,l}(t) \le x_i^j(t) \tag{16}$$

$$C13: \quad w_i^{j,l}(t) \ge x_i^j(t) - C(1 - z_i^l(t)) \tag{17}$$

$$C14: \quad w_i^{j,l}(t) \ge 0. \tag{18}$$

The exactness of this linearization follows directly. When $z_i^l(t) = 1$, constraints $C11$–$C13$ force $w_i^{j,l}(t) = x_i^j(t)$, whereas when $z_i^l(t) = 0$, constraints $C11$ and $C14$ imply $w_i^{j,l}(t) = 0$. Since $z_i^l(t)$ is binary and $x_i^j(t)$ is bounded in $[0, C]$, no relaxation gap is introduced. The reformed constraints become

$$C2': \quad Q_i^{j-1}(t+1) = Q_i^j(t) - \sum_{l\in\mathcal{L}} R_{i,l}(t)\, w_i^{j,l}(t) \tag{19}$$

$$C4': \quad \sum_{l\in\mathcal{L}} R_{i,l}(t)\, w_i^{j,l}(t) \le Q_i^j(t). \tag{20}$$

2) Linearizing the Objective Function: We replace the indicator in (3) with a binary surrogate $y_i \in \{0,1\}$, where $y_i = 1$ signifies that user i violates its delay-reliability target:

$$O_L: \quad \text{minimize} \sum_{i=1}^{N} y_i. \tag{21}$$

3) Linearizing Delay-Reliability Constraints: Per-slot violation indicators: Binary variables $\delta_i(t) \in \{0,1\}$ stand in for $\mathbb{I}(Q_i^0(t) > 0)$ and are enforced using big-M constraints:

$$C15: \quad Q_i^0(t) \le M_1 \cdot \delta_i(t), \quad \forall i, t \tag{22}$$

$$C16: \quad Q_i^0(t) \ge \epsilon \cdot \delta_i(t), \quad \forall i, t \tag{23}$$

with M_1 chosen large enough to upper-bound any feasible queue occupancy and $\epsilon > 0$ a small positive constant.

Violation ratio: Under the perfect-knowledge assumption, the total arrival count $A_i = \sum_{t=1}^{T} a_i(t)$ is known for each user. We pre-compute $\beta_i = 1/A_i$ and express

$$C17: \quad \gamma_i(T) = \beta_i \sum_{t=1}^{T} \delta_i(t), \quad \forall i. \tag{24}$$

4) Coupling Violation Ratio to User Satisfaction: The logical relation $\gamma_i(T) \ge \mu_i \Leftrightarrow y_i = 1$ is imposed through a second pair of big-M inequalities:

$$C18: \quad \gamma_i(T) - \mu_i \ge -M_2(1 - y_i), \quad \forall i \tag{25}$$

$$C19: \quad \gamma_i(T) - \mu_i \le M_2\, y_i, \quad \forall i, \tag{26}$$

where M_2 can be set to 1 because $\gamma_i(T)$ and μ_i both lie in $[0, 1]$.

5) Complete MILP: Collecting all components, the complete PODRX MILP is

$$\boxed{\begin{aligned} O_L: &\quad \text{minimize} \sum_{i=1}^{N} y_i \\ \text{subject to:} &\quad C1, C2', C4'\text{–}C10, \\ &\quad C11\text{–}C19, \\ &\quad y_i, \delta_i(t) \in \{0,1\}, \quad \forall i, t. \end{aligned}} \tag{27}$$

C. Optimal Solution

The MILP formulation in (27) can be solved to global optimality using standard solvers such as Gurobi or CPLEX. The main implementation subtlety arises from the recursive queue evolution in $C2'$, where each slot's queue state depends on the previous slot's state and decisions. Since these solvers do not natively handle recursive definitions, we unroll this recurrence over the entire horizon and express every queue variable $Q_i^j(t)$ as an explicit linear function of the decision variables and the known arrival sequence, then pass the resulting flattened constraint system to the solver.

Note that, PODRX is intended as a theoretical benchmark rather than a deployable scheduler. The MILP has exponential worst-case complexity and assumes perfect knowledge of future arrivals and channel states over the full scheduling horizon, which makes it unsuitable for real-time operation. Its role is to establish a performance ceiling: by quantifying the maximum gains achievable through joint power–RB optimization within the PDTQ framework, PODRX provides a theoretical upper bound against which practical causal algorithms, to be developed in future work, can be evaluated.

IV. Performance Evaluation

A. Simulation Setup

We evaluate PODRX using a 3GPP-compliant, custom system-level simulation framework implemented in Python. The scenario consists of a single gNB serving multiple XR users, with system parameters listed in Table II. Users are uniformly distributed over the cell coverage area. Each configuration is simulated over 20 independent runs with different random seeds, and results are averaged across these trials.

XR traffic is modeled according to the 3GPP framework in [13]. XR video traffic comprises pseudo-periodic frames, where PDU Set sizes and arrival jitter are drawn from truncated Gaussian distributions. For benchmarking, we compare PODRX with three standard schedulers, namely *Proportional Fair*

(PF), *Round Robin* (RR), and *Maximum Rate* (MR), as well as with the delay-reliability-aware DROX scheduler [10], which performs RB-only optimization at a fixed transmit power.

Performance is assessed using two key metrics. The primary metric is the *Satisfied XR User Count*, defined as the number of users for which at least 99% of PDU Sets are delivered within their PSDB, consistent with the 3GPP definition of XR capacity [5]. In addition, we report *Total System Power*, defined as the sum of the transmit power allocated across all users per slot, averaged over the simulation, to quantify the energy efficiency gains achieved through joint power–RB optimization.

Table II
SIMULATION PARAMETERS

Parameter	Value	Parameter	Value
Carrier freq.	3.5 GHz	Channel model	3GPP 38.901 [14]
BS total tx power per slot	43 dBm (20W)	UE tx power	23 dBm
Data rate (d_r)	30 Mbps	Frame rate (f_r)	60 fps
Numerology (μ)	1	Max bandwidth	100 MHz
PSDB	2.5–10 ms	Target delay-reliability (ρ)	99%
Mobility model	Stationary users	Mean PDU Set size ($\bar{S}$)	$\frac{d_r \times 10^6}{f_r \times 8}$ [Bytes]
Std. dev. PDU Set size (σ_S)	$0.105 \times \bar{S}$	PDU Set truncation range	$[0.5 \times \bar{S}, 1.5 \times \bar{S}]$
Mean inter-arrival time ($\bar{T}$)	$\frac{1}{f_r}$ [s]	Jitter mean ($\bar{J}$)	0 ms
Jitter std. dev. (σ_J)	2 ms	Cell range (R)	500m

B. Results and Discussions

We now examine the impact of key system parameters on the performance of the proposed PODRX scheduling scheme in comparison with existing approaches.

1) Effect of Number of XR Users: Fig. 3 compares scheduler performance under varying user loads. The baseline schedulers (PF, RR, MR) fail to satisfy more than one user at any load, which confirms that traditional scheduling is fundamentally inadequate for XR traffic. DROX achieves substantially better performance because it exploits the delay tracking mechanism, thereby confirming that accurate deadline awareness is a necessary ingredient for XR scheduling. PODRX consistently outperforms DROX across all user loads, and the performance gap widens as the load increases, growing from a marginal improvement at 5 users to approximately 20% more satisfied users at 20 users. This trend indicates that joint power–RB optimization becomes increasingly beneficial under stronger resource contention. Importantly, PODRX attains these gains while consuming more than 50% less power than DROX. This dual advantage arises because PODRX adapts transmit power based on both channel conditions and deadline urgency, avoiding the inefficiencies inherent in uniform power allocation. Notably, PODRX's power consumption increases with user count, reflecting the higher aggregate power required to serve

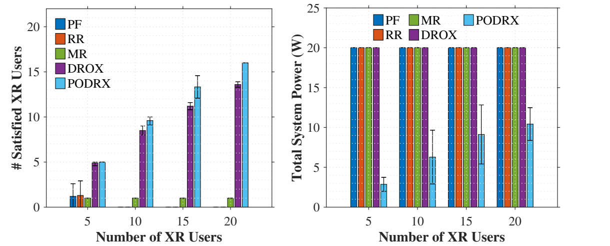

Figure 3. Effect of the user count: d_r = 30 Mbps; f_r = 60 fps; PSDB = 5 ms; ρ = 99%; f_c = 3.5 GHz; BW = 100 MHz; μ = 1; R = 500m.

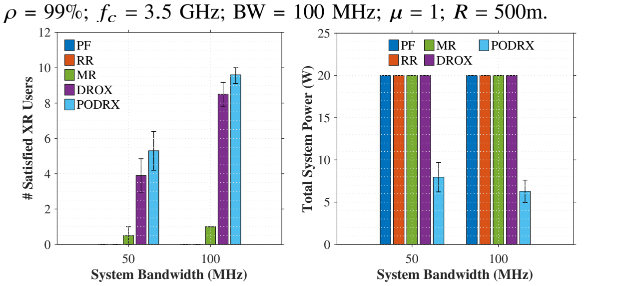

Figure 4. Effect of system bandwidth: N = 10; d_r = 30Mbps; f_r = 60fps; PSDB = 5ms; ρ = 99%; f_c = 3.5 GHz; μ = 1; R = 500m; Mobility: Stationary.

additional users, yet it remains well below the fixed power consumed by DROX regardless of load.

2) Effect of System Bandwidth: Fig. 4 reports performance as a function of system bandwidth for a fixed population of 10 XR users. Both DROX and PODRX benefit from additional bandwidth, yet PODRX consistently achieves higher user satisfaction and lower power consumption across all considered bandwidth values. The performance gap between PODRX and DROX decreases as bandwidth grows, from a 33% improvement at 50 MHz to 14% at 100 MHz, which reflects the diminishing marginal benefit of joint power–RB optimization when spectral resources are plentiful. A notable observation is that PODRX's total power consumption decreases with increasing bandwidth, since the availability of additional RBs reduces the need to operate at high power levels. This bandwidth–power trade-off illustrates PODRX's ability to exploit the substitutability between power and spectrum resources, making it particularly appealing in bandwidth-constrained deployments.

3) Effect of PDU Set Delay Budget: Fig. 5 analyzes the sensitivity to PSDB for a fixed set of 10 XR users. The performance of both DROX and PODRX improves as the delay budget is relaxed. Across the entire PSDB range, however, PODRX consistently maintains a 16–33% higher satisfied user count than DROX. The power consumption results exhibit a counterintuitive trend: PODRX consumes *less* power under tighter deadlines, achieving more than 90% power savings at PSDB = 2.5 ms. This behavior arises because stringent delay budgets render many users infeasible regardless of how power is allocated. PODRX effectively identifies and prioritizes only those users that can still meet their deadlines, rather than expending power on infeasible PDU Sets. This adaptive selection demonstrates that PODRX is particularly well suited to latency-critical applications in which energy efficiency is a

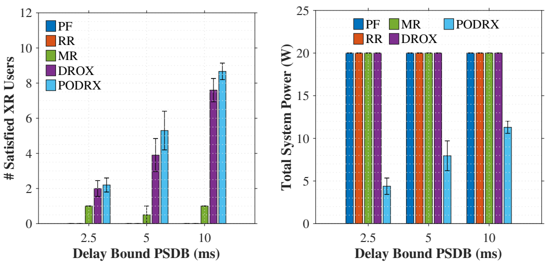

Figure 5. Effect of PSDB: d_r = 30Mbps; f_r = 60fps; ρ = 99%; f_c = 3.5 GHz; BW = 50MHz; μ = 1; R = 500m; Mobility: Stationary.

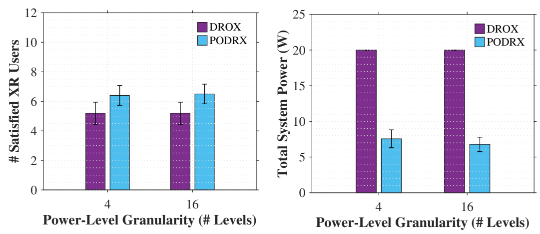

Figure 6. Effect of Power level granularity: d_r = 30Mbps; f_r = 60fps; ρ = 99%; f_c = 3.5 GHz; BW = 50MHz; PSDB = 5ms; μ = 1; R = 500m.

primary concern.

4) Effect of Power-Level Granularity: Fig. 6 evaluates the impact of power-level granularity by comparing 4-level and 16-level configurations. Here, granularity refers to the number of discrete transmit power options available to the optimizer, where 4 levels correspond to approximately 6 dB steps and 16 levels to approximately 1.5 dB steps between minimum and maximum power. PODRX achieves approximately 20% more satisfied users than DROX in both cases, confirming that user satisfaction gains are robust to the choice of power-level resolution. Power consumption, however, decreases with finer granularity, because finer resolution enables more precise power matching and reduces the overprovisioning inherent in coarser configurations. This benefit comes at the cost of increased computational complexity due to the larger search space of the decision variables.

V. Conclusion and Future Work

This paper has presented PODRX, a joint power and resource block allocation framework for delay-reliability-aware XR scheduling in 6G networks. By treating transmit power as a decision variable alongside RB allocation within the DTQ-based scheduling framework, PODRX exploits the coupling between power and achievable rate to improve both user satisfaction and energy efficiency. The scheduling problem was formulated as a MINLP and transformed into a tractable MILP, enabling globally optimal solutions that serve as performance benchmarks. Simulation results confirmed that PODRX consistently outperforms fixed-power schedulers across diverse operating conditions. The benefits were most pronounced in resource-constrained scenarios, such as tight delay budgets, limited bandwidth, or high user density, where intelligent power adaptation provides the greatest gains.

Future work will focus on developing low-complexity heuristics suitable for real-time implementation and on extending the framework to handle channel and traffic uncertainty.

Acknowledgments

This work was funded by Taighde Éireann-Research Ireland via grants 13/RC/2094_2 (Lero) and 21/FFP-P/10244 (FRADIS) and co-funded by the European Union (SyMeCo, Grant No. 101081459). The views expressed are solely those of the author(s). Neither the European Union nor the European Research Executive Agency can be held responsible for them.

References

[1] C.-X. Wang, X. You, X. Gao, X. Zhu, Z. Li, C. Zhang, H. Wang, Y. Huang, Y. Chen, H. Haas, J. S. Thompson, E. G. Larsson, M. D. Renzo, W. Tong, P. Zhu, X. Shen, H. V. Poor, and L. Hanzo, "On the road to 6g: Visions, requirements, key technologies, and testbeds," *IEEE Communications Surveys Tutorials*, vol. 25, no. 2, pp. 905–974, 2023.

[2] S. Gül, S. Ahsan, S. Paris, and I. D. D. Curcio, "Pdu set-based qos handling in 3gpp: Release 18 overview and future directions," *IEEE Communications Standards Magazine*, vol. 9, no. 3, pp. 104–110, 2025.

[3] M. Gapeyenko, S. Paris, M. Isomaki, B. Yanakiev, A. Amiri, B. Sébire, J. Kaikkonen, C. Wu, and K. I. Pedersen, "Overview of nr enhancements for extended reality (xr) in 3gpp 5g-advanced," *IEEE Access*, vol. 12, pp. 186 941–186 956, 2024.

[4] P. Paymard, S. Paris, A. Amiri, T. E. Kolding, F. S. Moya, and K. I. Pedersen, "Pdu-set scheduling algorithm for xr traffic in multi-service 5g-advanced networks," in *ICC 2024 - IEEE International Conference on Communications*, 2024, pp. 758–763.

[5] 3rd Generation Partnership Project (3GPP), "Study on XR (Extended Reality) Evaluations for NR (Release 17)," 3GPP, 3GPP Technical Report TR 38.838, Dec. 2021, available at: https://www.3gpp.org/ftp/Specs/archive/38_series/38.838/.

[6] X. Zhao, Y.-J. A. Zhang, M. Wang, X. Chen, and Y. Li, "Online multi-user scheduling for xr transmissions with hard-latency constraint: Performance analysis and practical design," *IEEE Transactions on Communications*, vol. 72, no. 7, pp. 4055–4071, 2024.

[7] E. Chen, S. Dou, S. Wang, Y. Cao, and S. Liao, "Frame-level integrated transmission for extended reality over 5g and beyond," in *2021 IEEE Global Communications Conference (GLOBECOM)*, 2021, pp. 1–6.

[8] S. Dutta, D. Roy, M. Laha, and G. Das, "Application-aware mac scheduling for xr over epon-based 6g-backhaul," *Optical Fiber Technology*, vol. 98, p. 104545, 2026. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1068520025004201

[9] M. Laha, D. Roy, S. Dutta, and G. Das, "Ai-assisted improved service provisioning for low-latency xr over 5g nr," *IEEE Networking Letters*, vol. 6, no. 1, pp. 31–35, 2024.

[10] M. Laha, G. Das, and G.-M. Muntean, "Delay-reliability aware optimal downlink scheduling for extended reality applications in 6g," in *IEEE INFOCOM 2025 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2025, pp. 1–6.

[11] M. Laha, G. Das, and G. Muntean, "Delay-reliability aware downlink scheduling for multi-sensory immersive extended reality in nextg," *IEEE Journal on Selected Areas in Communications*, vol. 44, pp. 3427–3443, 2026.

[12] A. Cedergren, F. de Laval, S. Falahati, L. Falk, D. H. Kang, R. S. Karlsson, Y. Lyazidi, A. Mirzaei, J. Olsson, J. L. Pradas, P. Schliwa-Bertling, N. Shi, B. Singh, R. Tano, and A. M. Voicu, "5g networks evolution for extended reality," *IEEE Communications Standards Magazine*, vol. 8, no. 3, pp. 54–59, 2024.

[13] "Traffic models and quality evaluation methods for media and xr services in 5g systems," in *3GPP TR 26.926 V0.1.0*, Apr. 2021.

[14] "Study on channel model for frequencies from 0.5 to 100 ghz," in *(3GPP), TR 38.901, V16.1.0*, Jan.2020.